

Archiver Storage

Notes from the Extremely Large
Database Conference

Eric Berryman
National Superconducting Cyclotron
Laboratory

Outline

- XLDB4
- Hadoop
- SciDB

- Extremely Large Databases
- 4th XLDB (first "open" conference)
October 6-7, 2010
SLAC
- <http://www-conf.slac.stanford.edu/xldb/>
- 158 Participants (with some neat industries)
 - Facebook
 - Yahoo!
 - Google
 - Cloudera

- Yahoo! and Cloudera and Facebook
 - Large community
- MapReduce
- Store data in Hadoop Distributed File System (HDFS)
 - Search files with Hive
- Or databases
 - HBase
 - Hypertable
- MapReduce: A major step backwards
 - <http://databasecolumn.vertica.com/database-innovation/mapreduce-a-major-step-backwards/>

- Started Oct 2007 resulting from the first XLBD conference
- <http://www.scidb.org/>
- Led by Mike Stonebraker - MIT
- Advisory board
 - Jacek Becla - SLAC
 - Kian-Tat Lim - SLAC
- Open Source
- Main sponsor is Zetics, but owner is scidb.org (non-profit)

SciDB - Why

- Science people are unhappy with RDBMS
 - Main data types are arrays (inefficient in tables)
 - Three main features missing
 - provenance (lineage)
 - uncertainty
 - version control
 - SQL operations don't fit (need array operations, ie. regrid)

SciDB - What is it?

- Shared nothing cluster parallelism 10's–1000's of nodes
- Array Oriented Data Model
- Append-only storage w/ support for
 - named versions
 - lineage (what generated the data)
 - time travel (don't overwrite, keep everything)
- Support for User Defined Functions (UDFs)
- Massively Parallel Computations
- Simple continuous model of uncertainty

- Multi-dimensional arrays
 - Integer-indexed dimensions
 - Cells contain scalars, User Defined Types (UDTs), or another array
- Coordinate systems map from user-defined types to integer indices in enhanced arrays
- Ragged arrays allow each row/column to have a different dimensionality
- User-definable handling for 'null' or missing data

- Array/analytics Query Language operation examples
 - Aggregate
 - Apply
 - Compose
 - Filter
 - Join (Combine)
 - Lookup
 - Project
 - Regid
 - Subsample
 - Multiple
 - Transpose
 - Or extend with UDFs

- Currently at R0.5
 - download link should be up now (if not, I can get you a copy)
 - pre-beta (Proof of Concept)
 - still working on developer process for community
- R0.75 end of the year
 - AQL
 - Error handling
 - Scalable math operations
 - Better documentation
- R1.0 April 2011
 - More functionally complete (UDFs, uncertainty, provenance)
 - Robust high performance

Thank you!